arXiv:1604.07319v1 [cs.CV] 25 Apr 2016

Semi-supervised Dictionary Learning Based on Hilbert-Schmidt Independence Criterion

Mehrdad J. Gangeh[1], Safaa M.A. Bedawi[2], Ali Ghodsi[3], and Fakhri Karray[2]

[1] Departments of Medical Biophysics, and Radiation Oncology,
University of Toronto, Canada
mehrdad.gangeh@utoronto.ca
[2] Center for Pattern Analysis and Machine Intelligence,
Department of Electrical and Computer Engineering, University of Waterloo, Canada
{sbedawi,karray}@uwaterloo.ca
[3] Department of Statistics and Actuarial Science, University of Waterloo, Canada
aghodsib@uwaterloo.ca

Abstract. In this paper, a novel semi-supervised dictionary learning and sparse representation (SS-DLSR) is proposed. The proposed method benefits from the supervisory information by learning the dictionary in a space where the dependency between the data and class labels is maximized. This maximization is performed using Hilbert-Schmidt independence criterion (HSIC). On the other hand, the global distribution of the underlying manifolds were learned from the unlabeled data by minimizing the distances between the unlabeled data and the corresponding nearest labeled data in the space of the dictionary learned. The proposed SS-DLSR algorithm has closed-form solutions for both the dictionary and sparse coefficients, and therefore does not have to learn the two iteratively and alternately as is common in the literature of the DLSR. This makes the solution for the proposed algorithm very fast. The experiments confirm the improvement in classification performance on benchmark datasets by including the information from both labeled and unlabeled data, particularly when there are many unlabeled data.

1 Introduction

Dictionary learning and sparse representation (DLSR) is one of the most successful mathematical models, which has led to state-of-the-art results in various applications such as face recognition [1,2,3], image denoising [4], texture classification [5], and emotion recognition [6]. DLSR, however, was originally proposed in an unsupervised setting [7]. The main objective function in the optimization problem related to DLSR is to minimize the reconstruction error between the original signal and the reconstructed one in the space of learned dictionary without including the information on class labels into the learning process. To formally describe the original DLSR formulation, we suppose that there is a finite set of data samples denoted as $\mathbf{X} = [\mathbf{x}_1, ..., \mathbf{x}_n] \in \mathbb{R}^{d \times n}$, where d is the dimensionality of the data and n is the number of data samples. In original

DLSR, the data is decomposed using a few dictionary atoms by optimizing the empirical cost function

$$L(\mathbf{X}, \mathbf{D}, \boldsymbol{\alpha}) = \sum_{i=1}^{n} l(\mathbf{x}_i, \mathbf{D}, \boldsymbol{\alpha}), \tag{1}$$

where $\mathbf{D} \in \mathbb{R}^{d \times k}$ is a dictionary of k atoms, $\boldsymbol{\alpha} \in \mathbb{R}^{k \times n}$ are the sparse coefficients and L, l are loss functions. In the literature of the DLSR, the reconstruction error, in mean-squared sense, between the original signal and the reconstructed signal is the most common loss function, which is usually regularized by the ℓ_1 norm to induce sparsity into the coefficients. Thus, the formulation in (1) can be written as

$$L(\mathbf{X}, \mathbf{D}, \boldsymbol{\alpha}) = \min_{\mathbf{D}, \boldsymbol{\alpha}} \sum_{i=1}^{n} \left(\tfrac{1}{2} \|\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i\|_2^2 + \lambda \|\boldsymbol{\alpha}_i\|_1 \right), \tag{2}$$

where $\boldsymbol{\alpha_i}$ is the ith column of α. In order to avoid arbitrarily large values for $\mathbf{D}$ and consequently, arbitrarily small values for α, we need an additional constraint on the dictionary atoms to limit their ℓ_2 norm to be smaller than or equal to one. The complete optimization problem in (2) after adding this constraint is as follows:

$$\begin{aligned} L(\mathbf{X}, \mathbf{D}, \boldsymbol{\alpha}) = \min_{\mathbf{D}, \boldsymbol{\alpha}} \quad & \sum_{i=1}^{n} \left(\tfrac{1}{2} \|\mathbf{x}_i - \mathbf{D}\boldsymbol{\alpha}_i\|_2^2 + \lambda \|\boldsymbol{\alpha}_i\|_1 \right), \\ \text{s.t.} \quad & \|\mathbf{d}_j\|_2^2 \leq 1 \quad \forall j = 1, .., , k. \end{aligned} \tag{3}$$

The original DLSR formulation given in (3) is unsupervised as the category information has not been taken into consideration in the optimization problem. However, in a supervised learning paradigm, where the ultimate goal is the classification of the data, this setting may not lead to an optimal discriminative dictionary nor coefficients. A more recent attempt in the literature was to incorporate the class labels into the learning of the dictionary and/or coefficients (refer to [8] for a review). This modification resulted in a new category of DLSR, namely called supervised dictionary learning and sparse representation (S-DLSR). Improvements (some significant) over unsupervised DLSR have been reported in the literature for the classification tasks [3,9,10,11].

Although S-DLSR benefits from the side information available from category information to learn a more discriminative dictionary, unfortunately, gathering labeled data is often very expensive and time consuming. Most data available is unlabeled and the sample size of the labeled data is often very small, which has a hindering effect on the discriminative quality of the learned dictionary. Semi-supervised learning (SSL) methods can potentially boost the performance of a machine learning system by utilizing both supervisory information and global data distribution. Using a large amount of unlabeled data, which is usually easily accessible, can improve revealing the manifold global distribution [12], and compensate for the small sample size of labeled data [13].

In this paper, a semi-supervised dictionary learning and sparse representation (SS-DLSR) based on Hilbert-Schmidt independence criterion (HSIC) is

proposed. The proposed SS-DLSR approach finds a dictionary based on two criteria: first, the maximization of the dependency between the labeled data and the corresponding category information, and second, minimization of the distances between the unlabeled data and their nearest labeled data. The first criterion guarantees finding the space of maximum discrimination based on the information in the category information and labeled data, whereas the second criterion, guarantees that the unlabeled data remain as close as possible to their nearest-neighbor labeled data. Therefore, the learned dictionary (the projection directions computed by using the aforementioned criteria) benefits from the discriminative power of the category information in the labeled data and proximity information of the unlabeled data as an indication of global manifold distribution. the sparse coefficients are subsequently computed in the space of learned dictionary using the formulation given in (3).

2 Semi-supervised Dictionary Learning and Sparse Representation

2.1 Problem Statement

Let $\mathbf{X} = [\mathbf{x}_1, ..., \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ be n data samples with the dimensionality of d. There are n_l labeled and n_u unlabeled data samples, where $n = n_l + n_u$. Let $\{(\mathbf{x}_1, \mathbf{y}_1), ..., (\mathbf{x}_{n_l}, \mathbf{y}_{n_l})\}$ be the pair of labeled data ($\mathbf{X}_l \in \mathbb{R}^{d \times n_l}$) and the corresponding labels ($\mathbf{Y} \in \{0, 1\}^{c \times n_l}$, where c is the number of classes), and $\mathbf{X}_u = [\mathbf{x}_{n_l+1}, ..., \mathbf{x}_n] \in \mathbb{R}^{d \times n_u}$ be the unlabeled data samples. We would like to find a dictionary, which can be considered as a transformation, based on two criteria (1) maximizing the dependency between the labeled data $\mathbf{X}_l$ and the labels $\mathbf{Y}$, and (2) minimizing the distance between each unlabeled data with the nearest label data. The first criterion is to guarantee finding a discriminative dictionary using the labeled data, and the second criterion is to ensure the unlabeled data samples are mapped close to their neighboring labeled data and therefore, the global connectivity of data is maintained in the space of the learned dictionary.

The first criterion is implemented using the Hilbert-Schmidt independence criterion (HSIC), which will be explained in the next subsection followed by the design of the dictionary and sparse coefficients for the proposed semi-supervised method.

2.2 Hilbert-Schmidt Independence Criterion

HSIC is a kernel-based measure of independence between two random variables $\mathcal{X}$ and $\mathcal{Y}$ proposed first by Gretton *et al.* [14,15]. It is computed based on the Hilbert-Schmidt norm of cross covariance operators in reproducing kernel Hilbert spaces (RKHSs) [15].

Our focus here is the empirical HSIC, which is computed using a finite set of data samples. To this end, considering $\mathcal{Z} := \{(\mathbf{x}_1, \mathbf{y}_1,), ..., (\mathbf{x}_{n_l}, \mathbf{y}_{n_l})\} \subseteq \mathcal{X} \times \mathcal{Y}$

as n_l independent observations drawn from joint probability distribution $P_{\mathcal{X}\times\mathcal{Y}}$, the empirical HSIC is computed using

$$\text{HSIC}(\mathcal{Z}) = \frac{1}{(n_l-1)^2}\text{tr}(\mathbf{KHBH}), \tag{4}$$

where tr is the trace operator, and $\mathbf{K}$, $\mathbf{B}$, $\mathbf{H} \in \mathbb{R}^{n_l \times n_l}$. $\mathbf{K}$ and $\mathbf{B}$ are kernels on the data and labels, respectively. $\mathbf{H} = \mathbf{I} - n_l^{-1}\mathbf{e}\mathbf{e}^\top$, where $\mathbf{I}$ is an identity matrix, $\mathbf{e}$ is a vector of all ones and therefore, $\mathbf{H}$ is a centering matrix. Since the empirical HSIC given in (4) is a measure of dependency between $\mathcal{X}$ and $\mathcal{Y}$, in order to maximize this dependency, tr($\mathbf{KHBH}$) should be maximized.

2.3 Dictionary Learning

As mentioned in the problem statement (Subsection 2.1), the dictionary is learned based on two criteria. In order to maximize the dependency between the labeled data and the corresponding labels, as shown in [11], the following optimization problem has to be solved:

$$\begin{aligned} \max_{\mathbf{D}} \quad & \text{tr}(\mathbf{D}^\top \mathbf{X}_l \mathbf{HBHX}_l^\top \mathbf{D}), \\ \text{s.t.} \quad & \mathbf{D}^\top \mathbf{D} = \mathbf{I} \end{aligned} \tag{5}$$

where $\mathbf{H}$ is the centering matrix, $\mathbf{B}$ is a kernel on labels, and $\mathbf{D}$ is the dictionary to be learned. By a few manipulations on the objective function given in (5), it can be demonstrated that it is another form of empirical HSIC:

$$\begin{aligned} &\max_{\mathbf{D}} \text{tr}(\mathbf{D}^\top \mathbf{X}_l \mathbf{HBHX}_l^\top \mathbf{D}) \\ &\quad = \max_{\mathbf{D}} \text{tr}(\mathbf{X}_l^\top \mathbf{D}\mathbf{D}^\top \mathbf{X}_l \mathbf{HBH}) \\ &\quad = \max_{\mathbf{D}} \text{tr}\left(\left[(\mathbf{D}^\top \mathbf{X}_l)^\top \mathbf{D}^\top \mathbf{X}_l\right]\mathbf{HBH}\right) \\ &\quad = \max_{\mathbf{D}} \text{tr}(\mathbf{KHBH}), \end{aligned} \tag{6}$$

where $\mathbf{K} = (\mathbf{D}^\top \mathbf{X}_l)^\top \mathbf{D}^\top \mathbf{X}_l$ is a linear kernel on the projected labeled data into the space of learned dictionary $\mathbf{D}$. As can be clearly observed from the last statement in (6), the objective function in (5) has the form of the empirical HSIC and thus, the dictionary $\mathbf{D}$ projects the labeled data to the space of maximum dependency with the corresponding labels.

The second criterion is to minimize the distances between the unlabeled data and the nearest neighbor labeled data in the space of the dictionary learned. In other words, considering $\mathbf{z} = \mathbf{D}^\top \mathbf{x}$ as a projected data sample to the space of the learned dictionary, we would like to:

$$\min_{\mathbf{D}} \frac{1}{2} \sum_{i=1}^{n_l} \sum_{j=1}^{n_u} w_{i,j} (\mathbf{z}_i - \mathbf{z}_j)^2, \tag{7}$$

where $w_{i,j}$ are the weights that define the proximity (neighborhood) of the unlabeled to labeled data. One way to define it is based one nearest neighbor, i.e., $w_{i,j} = 1$ if the jth unlabeled data is the nearest to the ith labeled data and $w_{i,j} = 0$ otherwise.

It can be shown [16] that the objective function given in (7) can be written in matrix form as follows:

$$\min_{\mathbf{D}} \frac{1}{2} \sum_{i=1}^{n_l} \sum_{j=1}^{n_u} w_{i,j} (\mathbf{z}_i - \mathbf{z}_j)^2 = \min_{\mathbf{D}} \quad \mathrm{tr}(\mathbf{Z}\mathbf{L}\mathbf{Z}^\top) = \min_{\mathbf{D}} \quad \mathrm{tr}(\mathbf{D}^\top \mathbf{X}\mathbf{L}\mathbf{X}^\top \mathbf{D}), \tag{8}$$

where $\mathbf{L}$ is the Laplacian of the graph made by the projected data points $\mathbf{Z} = [\mathbf{z}_1, ..., \mathbf{z}_n]$ in the space of learned dictionary, and is defined as $\mathbf{L} = \mathbf{Q} - \mathbf{W}$, where $\mathbf{W}(i,j) = w_{i,j}$ and $\mathbf{Q}$ is a diagonal matrix, where $q_{i,i} = \sum_j w_{i,j}$.

Combining the two objective functions given in (5) and 8, the overall optimization problem for the computation of the dictionary can be written as follows:

$$\begin{aligned} \max_{\mathbf{D}} \quad & \mathrm{tr}\left[\mathbf{D}^\top \left((1-\eta)\mathbf{X}_l \mathbf{H}\mathbf{B}\mathbf{H}\mathbf{X}_l^\top - \eta\, \mathbf{X}\mathbf{L}\mathbf{X}^\top\right)\mathbf{D}\right], \\ \text{s.t.} \quad & \mathbf{D}^\top \mathbf{D} = \mathbf{I} \end{aligned} \tag{9}$$

where $0 \leq \eta \leq 1$ is a constant that determines the relative contributions of the two terms in the objective function. According to the Rayleigh-Ritz theorem [17], the solution for the optimization problem given in (9) is the corresponding eigenvectors of the largest eigenvalues of $\mathbf{\Phi} = (1-\eta)\mathbf{X}_l \mathbf{H}\mathbf{B}\mathbf{H}\mathbf{X}_l^\top - \eta\, \mathbf{X}\mathbf{L}\mathbf{X}^\top$.

2.4 Sparse Coefficients

After the computation of the dictionary using (9), the sparse coefficients can be computed using the formulation provided in (2), which is called *lasso* if the dictionary is known [18]. Although (2) can be solved using fast iterative methods, since the dictionary is orthogonal, as shown in [19,20], the sparse coefficients can be computed using soft-thresholding with the soft-thresholding operator $S_\lambda(.)$:

$$\alpha_{ij} = S_\lambda \left([\mathbf{D}^\top \mathbf{x}_i]_j\right), \tag{10}$$

where $\alpha_{i,j}$ is the (i,j)th element of $\boldsymbol{\alpha}$ and $S_\lambda(t)$ is defined as follows:

$$S_\lambda(t) = \begin{cases} t - 0.5\lambda & \text{if } t > 0.5\lambda \\ t + 0.5\lambda & \text{if } t < -0.5\lambda \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

3 Experiments and Results

To validate the proposed semi-supervised dictionary learning and sparse representation method (SS-DLSR), two benchmark datasets publicly available from

UCI machine learning repository[1] were used. The two datasets were the Sonar ($n = 208$, $d = 60$, and $c = 2$) and the Parkinsons ($n = 297$, $d = 13$, and $c = 2$) datasets.

The performance of the proposed SS-DLSR was evaluated for a fixed dictionary size ($k = 8$) and varying relative ratio of the labeled to unlabeled data $n_l/(n_l + n_u)$. To this end, 70% of the data was randomly selected as the training set and 30% as the test set. The training data was further divided to different ratios of labeled and unlabeled data as shown in Table 1 ($n_l/(n_l + n_u) = \{0.05, 0.1, 0.3, 0.5\}$). One nearest neighbor was used as the proximity measure between the unlabeled and labeled data to determine the matrix of weights in (7). The value of η for the computation of the dictionary in (9) was set to three different values, i.e., 0 (ignoring unlabeled data), 1 (ignoring labeled data), and η^* (the most discriminative dictionary corresponding to best classification performance). The sparse coefficients were computed for the labeled portion of the training data as well as for the test data. A support vector machine (SVM) with radial basis function (RBF) kernel was used for the classification of the data by submission of the sparse coefficients to the classifier as suggested in [21]. The SVM was tuned using 5-fold cross validation on the labeled portion of training data to find the optimal kernel width (γ^*) and optimal trade-off parameter (C^*). Subsequently, the SVM was trained on whole labeled data in the training set using the optimal γ^* and C^* values and tested on the test set. The experiments were repeated 10 times for different random split of the data to training and test sets. The performance is reported in terms of classifier accuracy (averaged over 10 runs) in Table 1.

From the results provided in Table 1, there are several immediate observations. First, by adding unlabeled data to the learning of the dictionary (the columns in Table 1 corresponding with η^*), the classification performance is increased, which means that the learned dictionary is more discriminative. This reveals that the proposed algorithm can effectively incorporate the information from both labeled and unlabeled data into the learning of the dictionary. Second, by decreasing the rate of labeled to unlabeled data ($n_l/(n_l + n_u)$), the gain in performance from adding unlabeled data is increased. In realistic settings, there usually exist many unlabeled data and only a small number of labeled data. The proposed SS-DLSR algorithm benefits more from the information provided by the unlabeled data in these situations as can be observed by comparing the column corresponding with η^* (including both the labeled and unlabeled data into the dictionary learning) and the column with $\eta = 0$ (including only labeled data into the dictionary learning).

4 Discussion and Conclusion

In this paper, a novel semi-supervised dictionary learning and sparse representation method was proposed. A discriminative dictionary was learned in the

[1] http://archive.ics.uci.edu/ml/

Table 1: The classification rate (%) of the proposed SS-DLSR algorithm on two benchmark datasets. The results were compared for various settings in the proposed algorithm including different relative contributions of the labeled and unlabeled data on dictionary learning (varying η), and different ratios of labeled to unlabeled data (varying $n_l/(n_l + n_u)$).

$\frac{n_l}{n_l+n_u}$	Sonar			Ionosphere		
	$\eta = 0$	$\eta = 1$	η^*	$\eta = 0$	$\eta = 1$	η^*
0.5	69.03	51.29	70.97	88.45	76.90	86.72
	±5.78	±6.62	±3.95	±2.82	±5.15	±3.90
0.3	66.45	51.94	68.71	85.34	74.83	85.69
	±7.40	±4.01	±8.54	±3.92	±4.16	±3.99
0.1	57.74	49.19	61.45	78.94	73.45	80.34
	±4.97	±5.55	±8.13	±5.50	±5.28	±6.81
0.05	53.55	50.97	55.65	72.07	68.45	74.66
	±5.98	±5.44	±8.20	±13.52	±13.28	±6.56

space of maximum dependency between the labeled data and class labels, where the connectivity of the data was maintained by minimizing the distances between the unlabeled data and the corresponding nearest labeled data. As can be seen from (9), the dictionary has a closed form solution. Also, by using soft-thresholding, the sparse coefficients can be computed using a closed-form solution as given in (10). The proposed SS-DLSR approach is, therefore, very fast. The effectiveness of the proposed method in learning from both supervisory information (based on labeled data) and graph connectivity information (based on unlabeled data) was demonstrated by experiments on two benchmark datasets from UCI machine learning repository.

Acknowledgment. The first author gratefully acknowledges the funding from the Natural Sciences and Engineering Research Council (NSERC) of Canada under Postdoctoral Fellowship (PDF-454649-2014).

References

1. Zhong, C., Sun, Z., Tan, T.: Robust 3D face recognition using learned visual codebook. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2007) 1–6
2. Wright, J., Yang, A.Y., Ganesh, A., Sastry, S.S., Ma, Y.: Robust face recognition via sparse representation. IEEE Transactions on Pattern Analysis and Machine Intelligence **31**(2) (Feb. 2009) 210–227
3. Yang, M., Zhang, L., Feng, X., Zhang, D.: Fisher discrimination dictionary learning for sparse representation. In: 13$^{\text{th}}$ IEEE International Conference on Computer Vision (ICCV). (2011) 543–550
4. Mairal, J., Elad, M., Sapiro, G.: Sparse representation for color image restoration. IEEE Transactions on Image Processing **17**(1) (Jan. 2008) 53–69

5. Gangeh, M.J., Ghodsi, A., Kamel, M.S.: Dictionary learning in texture classification. In: Proceedings of the 8th international conference on Image analysis and recognition - Volume Part I, Berlin, Heidelberg, Springer-Verlag (2011) 335–343
6. Gangeh, M.J., Fewzee, P., Ghodsi, A., Kamel, M.S., Karray, F.: Multiview supervised dictionary learning in speech emotion recognition. IEEE/ACM Transactions on Audio, Speech and Language Processing **22**(6) (Jun. 2014) 1056–1068
7. Elad, M.: Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing. Springer, New York (2010)
8. Gangeh, M.J., Farahat, A.K., Ghodsi, A., Kamel, M.S.: Supervised dictionary learning and sparse representation-a review. CoRR **abs/1502.05928** (2015)
9. Mairal, J., Bach, F., Ponce, J., Sapiro, G., Zisserman, A.: Supervised dictionary learning. In: Advances in Neural Information Processing Systems (NIPS). (2008) 1033–1040
10. Wright, J., Ma, Y., Mairal, J., Sapiro, G., Huang, T.S., Yan, S.: Sparse representation for computer vision and pattern recognition. Proceedings of the IEEE **98**(6) (June 2010) 1031–1044
11. Gangeh, M.J., Ghodsi, A., Kamel, M.S.: Kernelized supervised dictionary learning. IEEE Transactions on Signal Processing **61**(19) (Oct. 2013) 4753–4767
12. Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Schölkopf, B.: Learning with local and global consistency. In: Advances in Neural Information Processing Systems (NIPS). (2004) 321–328
13. Chapelle, O., Schölkopf, B.: Semi-supervised Learning. MIT Press, Cambridge (2006)
14. Gretton, A., Herbrich, R., Smola, A.J., Bousquet, O., Schölkopf, B.: Kernel methods for measuring independence. Journal of Machine Learning Research **6** (Dec. 2005) 2075–2129
15. Gretton, A., Bousquet, O., Smola, A., Schölkopf, B.: Measuring statistical dependence with hilbert-schmidt norms. In: Proceedings of the 16th international conference on Algorithmic Learning Theory (ALT). (2005) 63–77
16. von Luxburg, U.: A tutorial on spectral clustering. Statistics and Computing **17**(4) (2007) 395–416
17. Lütkepohl, H.: Handbook of Matrices. John Wiley & Sons (1996)
18. Tibshirani, R.: Regression shrinkage and selection via the *lasso*. Journal of the Royal Statistical Society, Series B **58**(1) (1996) 267–288
19. Donoho, D.L., Johnstone, I.M.: Adapting to unknown smoothness via wavelet shrinkage. Journal of the American Statistical Association **90**(432) (1995) 1200–1224
20. Friedman, J., Hastie, T., Hofling, H., Tibshirani, R.: Pathwise coordinate optimization. The Annals of Applied Statistics **1**(2) (2007) 302–332
21. Raina, R., Battle, A., Lee, H., Packer, B., Ng, A.Y.: Self-taught learning: transfer learning from unlabeled data. In: Proceedings of the 24th international conference on Machine learning (ICML). (2007) 759–766